

Lecture Notes on Practical Foundations for Programming Languages

Part 4: Dynamic Classification

JIAWEN LIU, JÉRÉMY THIBAUT, and SHAOWEI ZHU

These are lecture notes taken during Robert Harper’s course on foundations of programming languages at the Oregon Programming Language Summer School 2019. See also a PFPL supplement on exceptions at <http://www.cs.cmu.edu/rwh/pfpl/supplements/exceptions.pdf>.

1 MOTIVATION

In this lecture, the goal is to build an exception control mechanism that:

- does not rely on any convention about the kind of data that is associated with an exception
- that ensures integrity and confidentiality of data: no third party should be able to intercept an exception intended for a particular handler

This is implemented by using dynamically created classes of data, and associating capabilities to them capabilities representing integrity and confidentiality of the associated data.

2 STATIC CLASSIFICATION

Dynamic classifications is to oppose to static classification, i.e., types built in the form of finite sums. Indeed, in a programming language equipped with static classification, a type is computed by solving a type equation:

$$\mathcal{U} \cong [\text{num} \hookrightarrow \omega, \text{nil} \hookrightarrow 1, \text{cons} \hookrightarrow \mathcal{U} \times \mathcal{U}, \text{fun} \hookrightarrow \mathcal{U} \rightarrow \mathcal{U}, \dots]$$

where each summand is a class of values.

Dynamic dispatch can be seen as a way to work with the data typed as above. When we want to manipulate a value of this data type, we need to *unfold* it and then pattern-match on the type.

$$\text{case}(u, \text{num} \hookrightarrow \dots; \dots; \text{cons} \hookrightarrow \langle x_1, x_2 \rangle)$$

The case analysis reveals further information about the type of the data that makes downstream manipulations possible.

Datatypes are *abstract* (statically "new": even you write the same implementation twice, they are two different types) recursive *sums* (static). Hence, this is an issue of using *generative* datatypes vs *concrete* datatypes.

3 DYNAMIC CLASSIFICATION

Now, we want a *dynamically extensible* sum: $\tau_1 + \tau_2 + \dots$, generating new datatypes at runtime.

This idea of dynamically “new” connects closely to the notions of information-flow or knowledge-flow: a dynamic new class corresponds to an unguessable secret (or perfect encryption, by α -conversion), and one could control the flow of information by choosing who knows the secret.

We obtain abstract datatypes, with *induced capabilities* that express integrity and confidentiality, where

- the implementer is responsible for the integrity of data;

- the client is responsible for confidentiality.

In particular, this may be used to implement exceptions: “shared secrets” between a raiser and a particular handler. In this case,

- the raiser is the implementer responsible for the integrity of the exception
- handlers other than the designated handler cannot decode the secret, thus are unable to handle the exception while being well-typed

Let us assume that there already exists an exception control-flow mechanism in our language. We will not focus on implementing the exception data itself.

4 IMPLEMENTATION OF THE EXCEPTION DATATYPE

Idea. We are going to associate with each exception a new class. An exception is raised by specifying its class; and an exception is handled by dispatching on the class.

4.1 Statics

$$\begin{aligned}
 \langle \text{Type} \rangle \tau &::= \dots \mid \underline{\text{clsfd}} \\
 \langle \text{Value} \rangle v &::= \dots \mid \underline{\text{in}}[c](v) && : \underline{\text{clsfd}} \\
 \langle \text{Comp} \rangle m &::= \dots \mid \underline{\text{isin}}[c](v, x.m_1, m_2) && \simeq \tau \\
 & \quad \mid \underline{\text{new}}[\tau](c.m)
 \end{aligned}$$

Here, c represents a class, used as an indexing scheme to distinguish between different exceptions. The computation $\underline{\text{isin}}[c](v, x.m_1, m_2)$ type checks as τ if both $m_1 \simeq \tau$ and $m_2 \simeq \tau$, and $v : \underline{\text{clsfd}}$. Its semantics is that it checks if v is in c , and performs $[x/v]m_1$ in that case, or m_2 otherwise. Finally, $\underline{\text{new}}[\tau](c.m)$ declares a new class c in m .

We also update our typing judgments to include an environment Σ of class binders associated with types.

$$\begin{aligned}
 \Gamma \vdash_{\Sigma} v &: \tau \\
 \Gamma \vdash_{\Sigma} m &\simeq \tau
 \end{aligned}$$

This environment Σ declares the active classes in the current scope. Note that, in our formalism, c is a class binder that is *not* a variable. In particular, the order in Σ does not matter, it is allowed to weaken Σ , and we can perform *injective* renaming, but not substitute (the renaming must be injective to preserve disequality).

The typing rule for “new” follows:

$$\frac{\Gamma \vdash_{\Sigma, a \sim \tau} m \simeq \tau'}{\Gamma \vdash_{\Sigma} \underline{\text{new}}[\tau](a.m) \simeq \tau'}$$

4.2 Dynamics

We consider the following:

- states $v\Sigma\{m\}$: classes v that could be used in computation m ;
- transitions $v\Sigma\{m\} \mapsto v\Sigma'\{m'\}$: where the invariant $\Sigma \subseteq \Sigma'$ holds;
- final states $v\Sigma\{\underline{\text{ret}}(v)\}$

We introduce a static judgment, meaning that a state is well-formed:

$$v\Sigma\{m\}\text{OK} \triangleq \vdash_{\Sigma} m \dot{\sim} \tau \text{ for any } \tau$$

The semantics of `new` is:

$$v\Sigma\{\text{new}[\tau](c.m)\} \mapsto v\Sigma, c \sim \tau\{m\}.$$

It allocates new classes of data, for “encryption” purposes.

The semantics of `in` and `isin` are:

$$v\Sigma, c \sim \tau, c' \sim \tau' \{ \text{isin}[c](\text{in}[c'](v'), x.m_1, m_2) \mapsto \begin{cases} [v'/x]m_1 & \text{if } c = c' \\ m_2 & \text{otherwise} \end{cases} .$$

The first case corresponds to the case where one learns the equality, and is able to perform the computation. In the second case, nothing at all is learnt.

4.3 Class references

Finally, we introduce the idea of class references.

$$\begin{aligned} \tau &::= \dots \mid \underline{\text{cls}} && \text{the type of class references} \\ v &::= \underline{\text{cls}}[c] && \text{written as } \&c \\ m &::= \underline{\text{reveal}}(v; c.m) \end{aligned}$$

The additional typing rules are:

$$\frac{}{\Gamma \vdash_{\Sigma, c} \tau \ \&c : \tau \ \underline{\text{cls}}} \qquad \frac{\Gamma \vdash_{\Sigma} v : \tau \ \underline{\text{cls}} \quad \Gamma \vdash_{\Sigma, c \sim \tau} m \dot{\sim} \tau'}{\Gamma \vdash_{\Sigma} \underline{\text{reveal}}(v; c.m) \dot{\sim} \tau'}$$

Reveal has the following semantics:

$$v\Sigma, c \sim \tau \{ \underline{\text{reveal}}(\&c; c.m) \} \mapsto v\Sigma, c \sim \tau\{m\}$$

This corresponds to guessing by α -conversion.

We may now define two shortcuts, that behaves like “in” and “isin”, but for class references:

$$\begin{aligned} \underline{\text{inref}}(v; v') &\triangleq \underline{\text{reveal}}(v; c.\text{in}[c](v')) \dot{\sim} \underline{\text{clsfd}} \\ \underline{\text{isinref}}(v; v'; x.m_1; m_2) &\triangleq \underline{\text{reveal}}(v; c.\text{isin}[c](v', x.m_1, m_2)). \end{aligned}$$

4.4 Possession of a class reference

Possession of a class reference affords *full rights* to create and analyze classified values. One may associate to a type $\tau \ \underline{\text{cls}}$ the following tuple:

$$\tau \ \underline{\text{cls}} \approx \left\{ \begin{array}{l} \text{mk} : \quad \tau \rightarrow \underline{\text{clsfd}} \\ \text{mk}^{-1} : \quad \underline{\text{clsfd}} \rightarrow \tau \ \underline{\text{opt}} \end{array} \right\}$$

and to the corresponding value &a the pair:

$$\&a \mapsto \left\{ \begin{array}{l} \text{mk} \hookrightarrow \lambda[\tau](x : \text{in}[a](x)) \\ \text{mk}^{-1} \hookrightarrow \lambda[\text{clsfd}](x.\text{isin}[a](x, y.\text{ret}(1.y), \text{ret}(2.<>))) \end{array} \right\}$$

The following properties hold:

- the possessor of mk is held responsible for the integrity of the underlying value, and can enforce the value respects an invariant
- the possessor of mk^{-1} can compromise the confidentiality, and confirm that the invariant holds

5 AN APPLICATION IN ML

Let us introduce some concepts from the ML language.

- exception X of τ declares an exception X carrying values of type τ
- raise e where $e : \text{exn}$ (that is, clsfd) raises an exception
- e handle $X(x)$ declares a handler. However, one might as well write z instead of $X(x)$, because one doesn't need to have access to the name of the exception to handle it.

So what does exception X of τ actually mean?

$$\text{val } X = \left\{ \begin{array}{l} \text{raise}[\tau] \quad \hookrightarrow \text{clsfd} \rightarrow \tau \\ \text{handle}[\tau] \quad \hookrightarrow \tau \rightarrow (\text{clsfd} \rightarrow \tau) \rightarrow \tau \end{array} \right\}$$

The handler takes two parameters: the expression to evaluate, and a function to call if an exception is raised.

In ML, raise has the following meaning, where $X : \tau \rightarrow \text{clsfd}$ and $x : \tau$:

$$\text{raise } X(v) \hookrightarrow \text{bnd}(\text{comp}(X.\text{mk}(v)), y.\text{RAISE}(y), z.\text{RAISE}(z))$$

This computes the content of the exception, classify it, and then calls the system implementation for raising an exception.

The second RAISE is used if the inner computation raises an exception itself.

Similarly, handle has the following meaning:

$$m \text{ handle } X(x) \Rightarrow m' \hookrightarrow \text{bnd}(\text{comp}(m), x.\text{ret}(x), y.\text{bnd}(\text{comp}(x.\text{mk}^{-1}(y)), x.m', z.\text{RAISE}(z)))$$

If the execution of m succeeds, the value is immediately returned. Otherwise, the exception is decoded (mk^{-1}), then handled by m' . Again, one need to handle the case of an inner exception.

Finally, it is now possible to see exceptions as dynamically allocating a new class:

$$\text{exception } X \text{ of } \tau \hookrightarrow \text{new}[\tau](c.\langle \text{mk} \hookrightarrow \dots, \text{mk}^{-1} \hookrightarrow \dots \rangle)$$